

---

# Training Neural Networks with Dropout Stochastic Hessian-Free Optimization

---

Ryan Kiros

Department of Computing Science  
University of Alberta  
Edmonton, AB, Canada  
rkiros@ualberta.ca

## Abstract

Hessian-free (HF) optimization has been successfully used for training deep autoencoders and recurrent networks. HF uses the conjugate gradient algorithm to construct update directions through curvature-vector products that can be computed on the same order of time as gradients. In this paper we exploit this property and study stochastic HF with small gradient and curvature mini-batches independent of the dataset size for classification. We modify Martens' HF for this setting and integrate dropout, a method for preventing co-adaptation of feature detectors, to guard against overfitting. On classification tasks, dropout stochastic HF achieves accelerated training and competitive results in comparison with dropout SGD without the need to tune learning rates.

## 1 Introduction

Stochastic gradient descent (SGD) has become the most popular algorithm for training neural networks. Not only is SGD simple to implement but its noisy updates often leads to solutions that are well-adapt to generalization on held-out data [1]. Furthermore, SGD operates on small mini-batches potentially allowing for scalable training on large datasets. For training deep networks, SGD can be used for fine-tuning after layerwise pre-training [2] which overcomes many of the difficulties of training deep networks. Additionally, SGD can be augmented with dropout [3] as a means of preventing overfitting.

There has been recent interest in second-order methods for training deep networks, partially due to the successful adaptation of Hessian-free (HF) by [4], an instance of the more general family of truncated Newton methods. Second-order methods operate in batch settings with less but more substantial weight updates. Furthermore, computing gradients and curvature information on large batches can easily be distributed across several machines. Martens' HF was able to successfully train deep autoencoders without the use of pre-training and was later used for solving several pathological tasks in recurrent networks [5].

HF iteratively proposes update directions using the conjugate gradient algorithm, requiring only curvature-vector products and not an explicit computation of the curvature matrix. Curvature-vector products can be computed on the same order of time as it takes to compute gradients with an additional forward and backward pass through the function's computational graph [6, 7]. In this paper we exploit this property and introduce stochastic HF, a variation of HF that operates on small gradient and curvature mini-batches independent of the dataset size. Our goal in developing stochastic HF is to combine the generalization advantages of SGD with second-order information from HF. Additionally we integrate dropout, as a means of preventing co-adaptation of feature detectors. We perform experimental evaluation on three datasets for classification: MNIST, USPS and Reuters,

obtaining accelerated training and competitive results in comparison with dropout SGD without the need to tune learning rates.

## 2 Related work

Much research has been investigated into developing adaptive learning rates or incorporating second-order information into SGD. [8] proposed augmenting SGD with a diagonal approximation of the Hessian while Adagrad [9] uses a global learning rate while dividing by the norm of previous gradients in its update. SGD with Adagrad was shown to be beneficial in training deep distributed networks for speech and object recognition [10]. To completely avoid tuning learning rates, [11] considered computing rates as to minimize estimates of the expectation of the loss at any one time. [12] proposed SGD-QN for incorporating a quasi-Newton approximation to the Hessian into SGD and used this to win one of the 2008 PASCAL large scale learning challenge tracks. Recently, [13] provided a relationship between HF, Krylov subspace descent and natural gradient due to their use of the Gauss-Newton curvature matrix. Furthermore, [13] argue that natural gradient is robust to overfitting as well as the order of the training samples. Other methods incorporating the natural gradient such as TONGA [14] have also showed promise on speeding up neural network training.

Analyzing the difficulty of training deep networks was done by [15], proposing a weight initialization that demonstrates faster convergence. More recently, [16] argue that large neural networks waste capacity in the sense that adding additional units fail to reduce underfitting on large datasets. The authors hypothesize the SGD is the culprit and suggest exploration with stochastic natural gradient or stochastic second-order methods. Such results further motivate our development of stochastic HF.

Related to our work is that of [17], who proposes a dynamic adjustment of gradient and curvature mini-batches for Hessian-free with convex losses based on variance estimations. Unlike our work, the batch sizes used are dynamic with a fixed ratio and are initialized as a function of the dataset size. Other work on using second-order methods for neural networks include [18] who proposed using the Jacobi pre-conditioner for Hessian-free, [19] using HF to generate text in recurrent networks and [20] who explored training with Krylov subspace descent (KSD). Unlike HF, KSD could be used with Hessian-vector products but requires additional memory to store a basis for the Krylov subspace. L-BFGS has also been successfully used in fine-tuning pre-trained deep autoencoders, convolutional networks [21] and training deep distributed networks [10]. Other developments and detailed discussion of gradient-based methods for neural networks is described in [22].

## 3 Hessian-free optimization

In this section we review Hessian-free optimization, largely following the implementation of Martens [4]. We refer the reader to [23] for detailed development and tips for using HF.

We consider unconstrained minimization of a function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  with respect to parameters  $\theta$ . More specifically, we assume  $f$  can be written as a composition  $f(\theta) = L(F(\theta))$  where  $L$  is a convex loss function and  $F(\theta)$  is the output of a neural network with  $\ell$  non-input layers. We will mostly focus on the case when  $f$  is non-convex. Typically  $L$  is chosen to be a matching loss to a corresponding transfer function  $p(z) = p(F(\theta))$ . For a single input, the  $(i+1)$ -th layer of the network is expressed as

$$y_{i+1} = s_i(W_i y_i + b_i) \quad (1)$$

where  $s_i$  is a transfer function,  $W_i$  is the weights connecting layers  $i$  and  $i+1$  and  $b_i$  is a bias vector. Common transfer functions include the sigmoid  $s_i(x) = (1 + \exp(-x))^{-1}$ , the hyperbolic tangent  $s_i(x) = \tanh(x)$  and rectified linear units  $s_i(x) = \max(x, 0)$ . In our work we strictly consider classification tasks, so the loss function used is the generalized cross entropy and softmax transfer

$$L(p(z), t) = -\sum_{j=1}^k t_j \log(p(z_j)), \quad p(z_j) = \exp(z_j) / \sum_{l=1}^k \exp(z_l) \quad (2)$$

where  $k$  is the number of classes,  $t$  is a target vector and  $z_j$  the  $j$ -th component of output vector  $z$ . Consider a local quadratic approximation  $M_\theta(\delta)$  of  $f$  around  $\theta$ :

$$f(\theta + \delta) \approx M_\theta(\delta) = f(\theta) + \nabla f(\theta)^T \delta + \frac{1}{2} \delta^T B \delta \quad (3)$$

where  $\nabla f(\theta)$  is the gradient of  $f$  and  $B$  is the Hessian or an approximation to the Hessian. If  $f$  was convex, then  $B \succeq 0$  and equation 3 exhibits a minimum  $\delta^*$ . In Newton's method,  $\theta_{k+1}$ , the parameters at iteration  $k + 1$ , are updated as  $\theta_{k+1} = \theta_k + \alpha_k \delta_k^*$  where  $\alpha_k \in [0, 1]$  is the rate and  $\delta_k^*$  is computed as

$$\delta_k^* = -B^{-1} \nabla f(\theta_{k-1}) \quad (4)$$

for which calculation requires  $O(n^3)$  time and thus often prohibitive. Hessian-free optimization alleviates this by using the conjugate gradient (CG) algorithm to compute an approximate minimizer  $\delta_k$ . Specifically, CG minimizes the quadratic objective  $q(\delta)$  given by

$$q(\delta) = \frac{1}{2} \delta^T B \delta + \nabla f(\theta_{k-1})^T \delta \quad (5)$$

for which the corresponding minimizer of  $q(\delta)$  is  $-B^{-1} \nabla f(\theta_{k-1})$ . The motivation for using CG is as follows: while computing  $B$  is expensive, compute the product  $Bv$  for some vector  $v$  can be computed on the same order of time as it takes to compute  $\nabla f(\theta_{k-1})$  using the R-operator [6]. Thus CG can efficiently compute an iterative solution to the linear system  $B\delta_k = -\nabla(f(\theta_{k-1}))$  corresponding to a new update direction  $\delta_k$ .

When  $f$  is non-convex, the Hessian may not be positive semi-definite and thus equation 3 no longer has a well defined minimum. Following Martens, we instead use the generalized Gauss-newton matrix defined as  $B = J^T L'' J$  where  $J$  is the Jacobian of  $f$  and  $L''$  is the Hessian of  $L$ . So long as  $f(\theta) = L(F(\theta))$  for convex  $L$  then  $B \succeq 0$ . Given a vector  $v$ , the product  $Bv = J^T L'' J v$  is computed successively by first computing  $Jv$ , then  $L''(Jv)$  and finally  $J^T(L''Jv)$  [7]. To compute  $Jv$ , we utilize the R-operator. The R-operator of  $F(\theta)$  with respect to  $v$  is defined as

$$\mathcal{R}_v\{F(\theta)\} = \lim_{\epsilon \rightarrow 0} \frac{F(\theta + \epsilon v) - F(\theta)}{\epsilon} = Jv \quad (6)$$

Computing  $\mathcal{R}_v\{F(\theta)\}$  in a neural network is easily done using a forward pass by computing  $\mathcal{R}_v\{y_i\}$  for each layer output  $y_i$ . More specifically,

$$\mathcal{R}_v\{y_{i+1}\} = \mathcal{R}_v\{W_i y_i + b_i\} s'_i = (v(W_i)y_i + v(b_i) + W_i \mathcal{R}\{y_i\}) s'_i \quad (7)$$

where  $v(W_i)$  is the components of  $v$  corresponding to parameters between layers  $i$  and  $i + 1$  and  $\mathcal{R}\{y_1\} = 0$  (where  $y_1$  is the input data). In order to compute  $J^T(L''Jv)$ , we simply apply back-propagation but using the vector  $L''Jv$  instead of  $\nabla L$  as is usually done to compute  $\nabla f$ . Thus,  $Bv$  may be computed through a forward and backward pass in the same sense that  $L$  and  $\nabla f = J^T \nabla L$  are.

As opposed to minimizing equation 3, Martens instead uses an additional damping parameter  $\lambda$  with damped quadratic approximation

$$\hat{M}_\theta(\delta) = f(\theta) + \nabla f(\theta)^T \delta + \frac{1}{2} \delta^T \hat{B} \delta = f(\theta) + \nabla f(\theta)^T \delta + \frac{1}{2} \delta^T (B + \lambda I) \delta \quad (8)$$

Damping the quadratic through  $\lambda$  gives a measure of how conservative the quadratic approximation is. A large value of  $\lambda$  is more conservative and as  $\lambda \rightarrow \infty$  updates become similar to stochastic gradient descent. Alternatively, a small  $\lambda$  allows for more substantial parameter updates especially along low curvature directions. Martens dynamically adjusts  $\lambda$  at each iteration using a Levenberg-Marquardt style update based on computing the reduction ratio

$$\rho = (f(\theta + \delta) - f(\theta)) / (M_\theta(\delta) - M_\theta(0)) \quad (9)$$

If  $\rho$  is sufficiently small or negative,  $\lambda$  is increased while if  $\rho$  is large then  $\lambda$  is decreased. The number of CG iterations used to compute  $\delta$  has a dramatic effect on  $\rho$  which is further discussed in section 4.1.

To accelerate CG, Martens makes use of the diagonal pre-conditioner

$$P = \left( \text{diag} \left( \sum_{j=1}^m \nabla f^{(j)}(\theta) \odot \nabla f^{(j)}(\theta) \right) + \lambda I \right)^{\xi} \quad (10)$$

where  $f^{(j)}(\theta)$  is the value of  $f$  for datapoint  $j$  and  $\odot$  denotes component-wise multiplication.  $P$  can be easily computed on the same backward pass as computing  $\nabla f$ .

Finally, two backtracking methods are used: one after optimizing CG to select  $\delta$  and the other a backtracking linesearch to compute the rate  $\alpha$ . Both these methods operate in the standard way, backtracking through proposals until the objective no longer decreases.

## 4 Stochastic Hessian-free

Martens' implementation utilizes the full dataset for computing objective values and gradients, and mini-batches for computing curvature-vector products. Naively setting both batch sizes to be small causes several problems. In this section we describe these problems and our contributions in modifying Martens' original algorithm to this setting.

### 4.1 Short CG runs, $\delta$ -momentum and use of mini-batches

The CG termination criteria used by Martens is based on a measure of relative progress in optimizing  $\hat{M}_\theta$ . Specifically, if  $x_j$  is the solution at CG iteration  $j$ , then training is terminated when

$$\frac{\hat{M}_\theta(x_j) - \hat{M}_\theta(x_{j-k})}{\hat{M}_\theta(x_j)} < \epsilon \quad (11)$$

where  $k = \max(10, j/10)$  and  $\epsilon$  is a small positive constant. The effect of this stopping criteria has a dependency on the strength of the damping parameter  $\lambda$ , among other attributes such as the current parameter settings. For sufficiently large  $\lambda$ , CG only requires 10-20 iterations when a preconditioner is used. As  $\lambda$  decreases, more iterations are required to account for pathological curvature that can occur in optimizing  $f$  and thus leads to more expensive CG iterations. Such behavior would be undesirable in a stochastic setting where preference would be put towards having equal length CG iterations throughout training. To account for this, we fix the number of CG iterations to be only 3-5 across training. Let  $\zeta$  denote this cut-off. Setting a limit on the number of CG iterations is used by [4] and [19] and also has a damping effect, since the objective function and quadratic approximation will tend to diverge as CG iterations increase [23]. We note that due to the short number of CG runs, the iterates from each solution are used during the CG backtracking step.

A contributor to the success of Martens' HF is the use of information sharing across iterations. At iteration  $k$ , CG is initialized to be the previous solution of CG from iteration  $k-1$ , with a small decay. For the rest of this work, we denote this as  $\delta$ -momentum.  $\delta$ -momentum helps correct proposed update directions when the quadratic approximation varies across iterations, in the same sense that momentum is used to share gradients. This momentum interpretation was first suggested by [23] in the context of adapting HF to a setting with short CG runs. Unfortunately, the use of  $\delta$ -momentum becomes challenging when short CG runs are used. Given a non-zero CG initialization,  $\hat{M}_\theta$  may be more likely to remain positive after terminating CG and assuming  $f(\theta + \delta) - f(\theta) < 0$ , means that the reduction ratio will be negative and thus  $\lambda$  will be increased to compensate. While this is not necessarily unwanted behavior, having this occur too frequently will push stochastic HF to behave more like SGD and possibly result in the backtracking linesearch to reject proposed updates. Our solution is to utilize a schedule on the amount of decay used on the CG starting solution. This is motivated by [23] suggesting more attention on the CG decay in the setting of using short CG runs. Specifically, if  $\delta_k^0$  is the initial solution to CG at iteration  $k$ , then

$$\delta_k^0 = \gamma_e \delta_{k-1}^\zeta, \quad \gamma_e = \min(1.01\gamma_{e-1}, .99) \quad (12)$$

where  $\gamma_e$  is the decay at epoch  $e$ ,  $\delta_1^0 = 0$  and  $\gamma_1 = 0.5$ . While in batch training a fixed  $\gamma$  is suitable, in a stochastic setting it is unlikely that a global decay parameter is sufficient. Our schedule has an annealing effect in the sense that  $\gamma$  values near 1 are feasible late in training even with only 3-5 CG

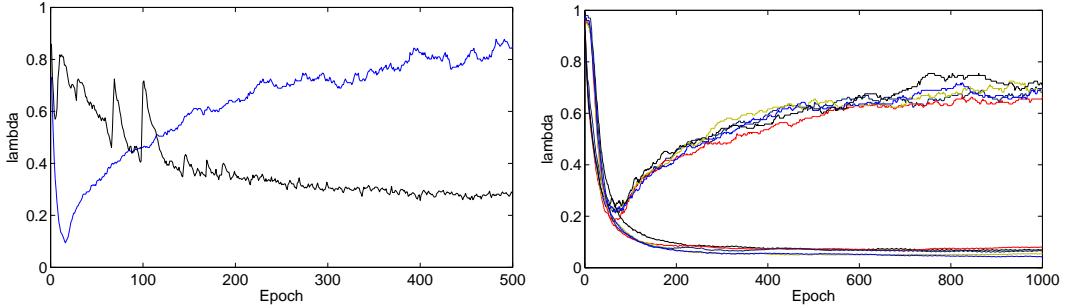


Figure 1: Values of the damping strength  $\lambda$  during training of MNIST (left) and USPS (right) with and without dropout using  $\lambda = 1$ . When dropout is included, the damping strength initially decreases followed by a steady increase over time.

iterations, a property that is otherwise hard to achieve. This allows us to benefit from sharing more information across iterations late in training, similar to that of a typical momentum method.

A remaining question to consider is how to set the sizes of the gradient and curvature mini-batches. [23] discuss theoretical advantages to utilizing the same mini-batches for computing the gradient and curvature vector products. In our setting, this leads to some difficulties. Using same-sized batches allows  $\lambda \rightarrow 0$  during training [23]. Unfortunately, this becomes incompatible with our short hard-limit on the number of CG iterations, since CG requires more work to optimize  $\hat{M}_\theta$  when  $\lambda$  approaches zero. To account for this, we opt to use gradient mini-batches that are 5-10 times larger than curvature mini-batches and in experiments use mini-batches of size 1000 and 100, respectively. In this setting, the behavior of  $\lambda$  is dependent on whether or not dropout is used during training. Figure 1 demonstrates the behavior of  $\lambda$  during training with and without the use of dropout. With dropout,  $\lambda$  no longer converges to 0 but instead plummets, rises and flattens out. In both settings,  $\lambda$  does not decrease substantially as to negatively effect the proposed CG solution and consequently the reduction ratio. Thus, the amount of work required by CG remains consistent late in training. The other benefit to using larger gradient batches is to account for the additional computation in computing curvature-vector products which would make training longer if both mini-batches were small and of the same size. In [4], the gradients and objectives are computed using the full training set throughout the algorithm, including during CG backtracking and the backtracking linesearch. We utilize the gradient mini-batch for the current iteration in order to compute all necessary gradient and objectives throughout the algorithm.

## 4.2 Levenberg-Marquardt damping

Martens makes use of the following Levenberg-Marquardt style damping criteria for updating  $\lambda$ :

$$\text{if } \rho > \frac{3}{4}, \lambda \leftarrow \frac{2}{3}\lambda \text{ elseif } \rho < \frac{1}{4}, \lambda \leftarrow \frac{3}{2}\lambda \quad (13)$$

which given a suitable initial value will converge to zero as training progresses. We observed that the above damping criteria is too harsh in the stochastic setting in the sense that  $\lambda$  will frequently oscillate, which is sensible given the size of the curvature mini-batches. We instead opt for a much softer criterion, for which lambda is updated as

$$\text{if } \rho > \frac{3}{4}, \lambda \leftarrow \frac{99}{100}\lambda \text{ elseif } \rho < \frac{1}{4}, \lambda \leftarrow \frac{100}{99}\lambda \quad (14)$$

This choice, although somewhat arbitrary, is consistently effective. Thus reduction ratio values computed from curvature mini-batches will have less overall influence on the damping strength.

## 4.3 Integrating dropout

Dropout is a recently proposed method for improving the training of neural networks. During training, each hidden unit is omitted with a probability of 0.5 along with optionally omitting input features similar to that of a denoising autoencoder [24]. Dropout can be viewed in two ways. By

randomly omitting feature detectors, dropout prevents co-adaptation among detectors which can improve generalization accuracy on held-out data. Secondly, dropout can be seen as a type of model averaging. At test time, outgoing weights are halved. If we consider a network with a single hidden layer and  $k$  feature detectors, using the mean network at test time corresponds to taking the geometric average of  $2^k$  networks with shared weights. Dropout is integrated in stochastic HF by randomly omitting feature detectors on both gradient and curvature mini-batches from the last hidden layer during each iteration. Since we assume that the curvature mini-batches are a subset of the gradient mini-batches, the same feature detectors are omitted in both cases.

Since the curvature estimates are noisy, it is important to consider the stability of updates when different stochastic networks are used in each computation. The weight updates in dropout SGD are augmented with momentum not only for stability but also to speed up learning. Specifically, at iteration  $k$  the parameter update is given by

$$\Delta\theta_k = p_k \Delta\theta_{k-1} - (1 - p_k) \alpha_k \langle \nabla f \rangle, \quad \theta_k = \theta_{k-1} + \Delta\theta_k \quad (15)$$

where  $p_k$  and  $\alpha_k$  are the momentum and learning rate, respectively. We incorporate an additional exponential decay term  $\beta_e$  when performing parameter updates. Specifically, each parameter update is computed as

$$\theta_k = \theta_{k-1} + \beta_e \alpha_k \delta_k, \quad \beta_e = c \beta_{e-1} \quad (16)$$

where  $c \in (0, 1]$  is a fixed parameter chosen by the user. Incorporating  $\beta_e$  into the updates, along with the use of  $\delta$ -momentum, leads to more stable updates and fine convergence when dropout is integrated during training.

#### 4.4 Algorithm

Pseudo-code for one iteration of our implementation of stochastic Hessian-free is presented. Given a gradient minibatch  $X_k^g$  and curvature minibatch  $X_k^c$ , we first sample dropout units for the inputs and last hidden layer of the network. These take the form of a binary vector, which are multiplied component-wise by the activations  $y_i$ . In our pseudo-code,  $\text{CG}(\delta_k^0, \nabla f, P, \zeta)$  is used to denote applying CG with initial solution  $\delta_k^0$ , gradient  $\nabla f$ , pre-conditioner  $P$  and  $\zeta$  iterations. Note that, when computing  $\delta$ -momentum, the  $\zeta$ -th solution in iteration  $k-1$  is used as opposed to the solution chosen via backtracking. Given the objectives  $f_{k-1}$  computed with  $\theta$  and  $f_k$  computed with  $\theta + \delta_k$ , the reduction ratio  $\rho$  is calculated utilizing the un-damped quadratic approximation  $M_\theta(\delta_k)$ . This allows updating  $\lambda$  using the Levenberg–Marquardt style damping. Finally, a backtracking linesearch with at most  $\omega$  steps is performed to compute the rate and serves as a last defense against potentially poor update directions.

Since curvature mini-batches are sampled from a subset of the gradient mini-batch, it is then sensible to utilize different curvature mini-batches on different epochs. Along with cycling through gradient mini-batches during each epoch, we also cycle through curvature subsets every  $h$  epochs, where  $h$  is the size of the gradient mini-batches divided by the size of the curvature mini-batches. Thus in our experiments, curvature mini-batch sampling completes a full cycle every  $1000/100 = 10$  epochs.

Finally, one simple way to speed up training as indicated in [23], is to cache the activations when initially computing the objective  $f_k$ . While each iteration of CG requires computing a curvature-vector product, the network parameters are fixed during CG and is thus wasteful to re-compute the network activations on each iteration.

## 5 Experiments

We perform classification experiments on three datasets: MNIST, USPS and Reuters. MNIST <sup>1</sup> is a collection of  $28 \times 28$  handwritten digits partitioned into 60000 digits for training and 10000 for testing. USPS <sup>2</sup> is a collection of 11000  $16 \times 16$  handwritten digits. We randomly construct 5 partitions of 8000 training and 3000 testing images. Reuters <sup>3</sup> [25] is a collection of 8293

---

<sup>1</sup><http://yann.lecun.com/exdb/mnist/>

<sup>2</sup><http://www.cs.nyu.edu/~roweis/data.html>

<sup>3</sup><http://www.cad.zju.edu.cn/home/dengcai/Data/TextData.html>

---

**Algorithm 1** Stochastic Hessian-Free Optimization

---

```

 $X_k^g \leftarrow$  gradient minibatch,  $X_k^c \leftarrow$  curvature minibatch,  $|X_k^g| = h|X_k^c|, h \in \mathbb{Z}^+$ 
Sample dropout units for inputs and last hidden layer
if start of new epoch then
     $\gamma_e \leftarrow \min(1.01\gamma_{e-1}, .99)$  { $\delta$ -momentum}
end if
 $\delta_k^0 \leftarrow \gamma_e \delta_{k-1}^\zeta$ 
 $f_{k-1} \leftarrow f(X_k^g; \theta), \nabla f \leftarrow \nabla f(X_k^g; \theta), P \leftarrow \text{Precon}(X_k^g; \theta)$ 
Solve  $(B + \lambda I)\delta_k = -\nabla f$  using CG( $\delta_k^0, \nabla f, P, \zeta$ ) {Using  $X_k^c$  to compute  $B\delta_k$ }
 $f_k \leftarrow f(X_k^g; \theta + \delta_k)$  {CG backtracking}
for  $j = \zeta - 1$  to 1 do
     $f(\theta + \delta_k^j) \leftarrow f(X_k^g; \theta + \delta_k^j)$ 
    if  $f(\theta + \delta_k^j) < f_k$  then
         $f_k \leftarrow f(\theta + \delta_k^j), \delta_k \leftarrow \delta_k^j$ 
    end if
end for
 $\rho \leftarrow (f_k - f_{k-1}) / (\frac{1}{2}\delta_k^T B\delta_k + \nabla f^T \delta_k)$  {Using  $X_k^c$  to compute  $B\delta_k$ }
if  $\rho < .25, \lambda \leftarrow 1.01\lambda$  elseif  $\rho > .75, \lambda \leftarrow .99\lambda$  end if {Backtracking linesearch}
 $\alpha_k \leftarrow 1, j \leftarrow 0$ 
while  $j < \omega$  do
    if  $f_k > f_{k-1} + .01\alpha_k \nabla f^T \delta_k$  then  $\alpha_k \leftarrow .8\alpha_k, j \leftarrow j + 1$  else break end if
end while
 $\theta \leftarrow \theta + \beta_e \alpha_k \delta_k, k \leftarrow k + 1$  {Parameter update}

```

---

documents labeled from 65 sub-classes. The publically available term frequencies and train/test split is used. Code for training SHF and reproducing our experimental results is available at <http://www.ualberta.ca/~rkiros/>.

## 5.1 Parameter settings

We train networks of size 784-1200-1200-10 for MNIST, 256-500-500-10 for USPS and 18900-65 for Reuters, for which the later reduces to softmax regression which is convex in  $f$ . Rectified linear units (ReLUs) are used for the activation function. For both SHF and dropout SGD, we tune only the percentage of input corruption. The training set is split in half which one chunk is used for testing. SHF and SGD are ran for a fixed number of epochs using corruption percentages of 20% and 50% from which the best result is used for training on the full training set. All experiments use a similar sparse initialization to Martens [4] with initial biases set to 0.1. The sparse initialization in combination with ReLUs makes our networks similar to the deep sparse rectifier networks of [26]. For dropout SHF, we initialize  $\lambda = 1$  and use a weight decay of  $2 \times 10^{-5}$ . Additionally, we experiment with SHF without the use of dropout and input corruption. In this setting, all experiments use a larger weight decay of  $5 \times 10^{-4}$ . No weight decay is used for SGD as in [3]. When training SGD we utilize the same weight clipping squared norm of 15 along with mini-batch sizes of 100 and the same momentum schedule as was done by [3], in conjunction with an initial learning rate of 10. The max squared norm allows SGD to use large initial learning rates for greater exploration early in training. Prior to training, MNIST and USPS are normalized to have zero mean and unit variance. Reuters is pre-processed by applying  $\log(1 + C)$  to word counts  $C$ . For additional comparison, we also train dropout SGD where dropout is only applied on the last hidden layer, as is done with dropout SHF.

## 5.2 Results

Figure 2 presents training and testing error curves for dropout SGD and SHF with and without dropout. On MNIST, dropout SHF demonstrates accelerated training, first breaking 120 errors after 50 epochs and 110 errors after 70 epochs. Both algorithms use 50% input corruption. [3] required a few hundred epochs to surpass 110 errors when logistic activations and 20% input corruption is used. At epoch 500, dropout stochastic HF achieves 107 errors. This result is similar to [3] which achieve

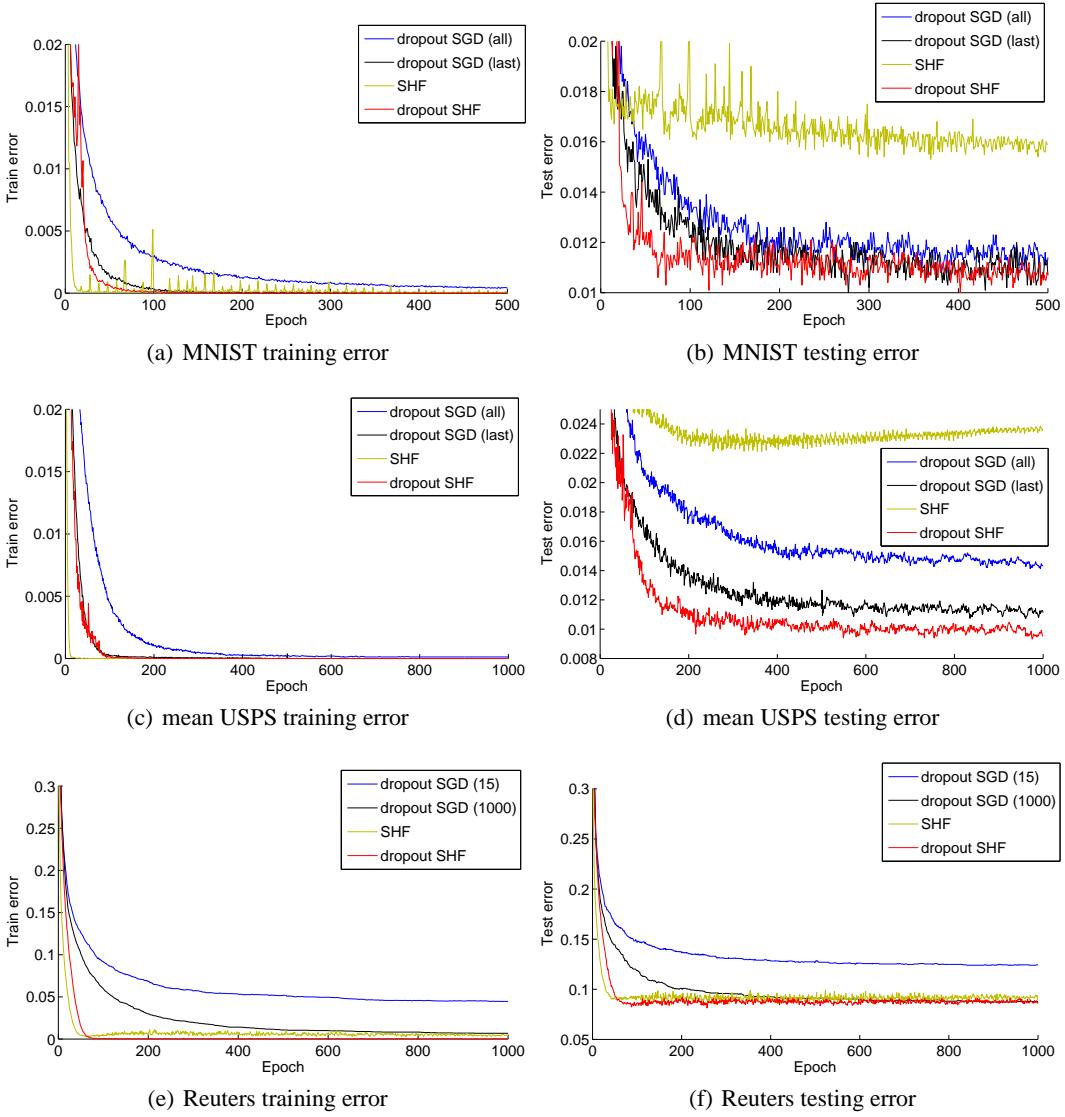


Figure 2: Results of applying dropout SGD and stochastic HF with and without dropout. For USPS, the mean of all 5 splits across epochs is presented. For MNIST and USPS, dropout SGD (all) refers to applying dropout to all hidden layers while dropout SGD (last) refers to only applying dropout on the last hidden layer. For Reuters, dropout SGD (15) refers to weight columns having a max squared norm of 15 and similarly for dropout SGD (1000).

100-115 errors with various network sizes when training for a few thousand epochs. Without dropout or input corruption, SHF achieves 159 errors on MNIST, on par with existing methods that do not incorporate prior knowledge, pre-training and image distortions. As with [4], we hypothesize that further improvements can be made by fine-tuning with stochastic HF after unsupervised layerwise pre-training.

After 1000 epochs of training on five random splits of USPS, we obtain final classification errors of 1%, 1.1%, 0.8%, 0.9% and 0.97% with a mean test error of 0.95%. Both algorithms use 50% input corruption. For additional comparison, [27] obtains a mean classification error of 1.14% using a pre-trained deep network for large-margin nearest neighbor classification with the same size splits. Interestingly, only using dropout in the last hidden layer improves performance. Without dropout, SHF overfits the training data.

On the Reuters dataset, SHF with and without dropout both demonstrate accelerated training. We hypothesize that further speedup may also be obtained by starting training with a much smaller  $\lambda$  initialization. Furthermore, max norm constraints may also be incorporated into SHF allowing for smaller  $\lambda$  initializations in the same sense as it allows SGD to begin training with larger learning rates.

In terms of the CPU training time, both dropout SGD and stochastic HF have similar performances per epoch given the batch sizes used. Using smaller gradient batches and/or larger curvature batches would increase the training time per epoch due to the extra forward and backward passes needed to computer curvature-vector products.

## 6 Conclusion

In this paper we proposed a stochastic variation of Martens' Hessian-free optimization incorporating dropout for training neural networks on classification problems. Our approach removes the need to tune learning rate and momentum schedules while obtaining competitive performance with dropout SGD on three datasets. While our initial results are promising, of interest would be adapting dropout stochastic HF to other network architectures:

- **Training deeper networks.** Our experimental evaluation consider networks with up to two hidden layers, which is often sufficient to obtain strong performance on most tasks. On the other hand, tasks such as acoustic modeling have largely benefit from using deeper networks [28].
- **Convolutional networks.** The most common approach to training convolutional networks has been SGD incorporating a diagonal Hessian approximation [8]. Dropout SGD was recently used for training a deep convolutional network on ImageNet [29].
- **Recurrent Networks.** It was largely believed that RNNs were too difficult to train with SGD due to the exploding/vanishing gradient problem. In recent years, recurrent networks have become popular again due to several advancements made in their training [30].
- **Recursive Networks.** Recursive networks have been successfully used for tasks such as sentiment classification and compositional modeling of natural language from word embeddings [31]. These architectures are usually trained using L-BFGS.

A downside to our approach is the number of heuristics used in setting and adapting parameters. It is not clear yet whether this setup is easily generalizable to the above architectures or whether improvements need to be considered. Furthermore, additional experimental comparison would involve dropout SGD with the adaptive methods of Adagrad [9] or [11], as well as the importance of pre-conditioning CG. None the less, we hope that this work initiates future research in developing stochastic Hessian-free algorithms.

## Acknowledgments

The author would like to thank Csaba Szepesvari for helpful discussion as well as David Sussillo for his guidance when first learning about and implementing HF. The author would also like to thank the anonymous ICLR reviewers for their comments and suggestions.

## References

- [1] L. Bottou and O. Bousquet. The tradeoffs of large-scale learning. *Optimization for Machine Learning*, page 351, 2011.
- [2] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy layer-wise training of deep networks. *NIPS*, 19:153, 2007.
- [3] G.E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R.R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv:1207.0580*, 2012.
- [4] J. Martens. Deep learning via hessian-free optimization. In *ICML*, volume 951, 2010.
- [5] J. Martens and I. Sutskever. Learning recurrent neural networks with hessian-free optimization. In *ICML*, 2011.

- [6] B.A. Pearlmutter. Fast exact multiplication by the hessian. *Neural Computation*, 6(1):147–160, 1994.
- [7] N.N. Schraudolph. Fast curvature matrix-vector products for second-order gradient descent. *Neural computation*, 14(7):1723–1738, 2002.
- [8] Y. LeCun, L. Bottou, G. Orr, and K. Müller. Efficient backprop. *Neural networks: Tricks of the trade*, pages 546–546, 1998.
- [9] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *JMLR*, 12:2121–2159, 2010.
- [10] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, Q. Le, M. Mao, A. Senior, P. Tucker, K. Yang, et al. Large scale distributed deep networks. In *NIPS*, pages 1232–1240, 2012.
- [11] T. Schaul, S. Zhang, and Y. LeCun. No more pesky learning rates. *arXiv:1206.1106*, 2012.
- [12] A. Bordes, L. Bottou, and P. Gallinari. Sgd-qn: Careful quasi-newton stochastic gradient descent. *JMLR*, 10:1737–1754.
- [13] Razvan Pascanu and Yoshua Bengio. Natural gradient revisited. *arXiv preprint arXiv:1301.3584*, 2013.
- [14] N. Le Roux, P.A. Manzagol, and Y. Bengio. Topmoumoute online natural gradient algorithm. In *NIPS*, 2007.
- [15] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 2010.
- [16] Yann N Dauphin and Yoshua Bengio. Big neural networks waste capacity. *arXiv preprint arXiv:1301.3583*, 2013.
- [17] R.H. Byrd, G.M. Chin, J. Nocedal, and Y. Wu. Sample size selection in optimization methods for machine learning. *Mathematical Programming*, pages 1–29, 2012.
- [18] O. Chapelle and D. Erhan. Improved preconditioner for hessian free optimization. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [19] I. Sutskever, J. Martens, and G. Hinton. Generating text with recurrent neural networks. In *ICML*, 2011.
- [20] O. Vinyals and D. Povey. Krylov subspace descent for deep learning. *arXiv:1111.4259*, 2011.
- [21] Q.V. Le, J. Ngiam, A. Coates, A. Lahiri, B. Prochnow, and A.Y. Ng. On optimization methods for deep learning. In *ICML*, 2011.
- [22] Y. Bengio. Practical recommendations for gradient-based training of deep architectures. *arXiv:1206.5533*, 2012.
- [23] J. Martens and I. Sutskever. Training deep and recurrent networks with hessian-free optimization. *Neural Networks: Tricks of the Trade*, pages 479–535, 2012.
- [24] P. Vincent, H. Larochelle, Y. Bengio, and P.A. Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, pages 1096–1103, 2008.
- [25] Deng Cai, Xuanhui Wang, and Xiaofei He. Probabilistic dyadic data analysis with local and global consistency. In *ICML*, pages 105–112, 2009.
- [26] X. Glorot, A. Bordes, and Y. Bengio. Deep sparse rectifier neural networks. In *AISTATS*, 2011.
- [27] R. Min, D.A. Stanley, Z. Yuan, A. Bonner, and Z. Zhang. A deep non-linear feature mapping for large-margin knn classification. In *ICDM*, pages 357–366, 2009.
- [28] A. Mohamed, G.E. Dahl, and G. Hinton. Acoustic modeling using deep belief networks. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(1):14–22, 2012.
- [29] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. *NIPS*, 25, 2012.
- [30] Y. Bengio, N. Boulanger-Lewandowski, and R. Pascanu. Advances in optimizing recurrent networks. *arXiv:1212.0901*, 2012.
- [31] R. Socher, B. Huval, C.D. Manning, and A.Y. Ng. Semantic compositionality through recursive matrix-vector spaces. In *EMNLP*, pages 1201–1211, 2012.